Novelty Assessment Report

Paper: Quantized Visual Geometry Grounded Transformer

PDF URL: https://openreview.net/pdf?id=Xzcllrc6gb

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-29

Abstract

Learning-based 3D reconstruction models, represented by Visual Geometry Grounded Transformers (VGGTs), have achieved remarkable progress with large-scale transformers. Their prohibitive computational and memory costs severely hinder real-world deployment. Post-Training Quantization (PTQ) has emerged as a common practice to compress and accelerate models. However, we empirically observe that PTQ faces unique obstacles when compressing billion-scale VGGTs: the data-independent special tokens induce heavy-tailed activation distributions, while the multi-view nature of 3D data makes calibration sample selection highly unstable. This paper proposes the first **Quant**ization framework for **VGGT**s, namely **QuantVGGT**. This mainly relies on two technical contributions: First, we introduce Dual-Smoothed Fine-Grained Quantization, which integrates pre-global Hadamard rotation and post-local channel smoothing to robustly mitigate heavy-tailed distributions and inter-channel variance. Second, we design Noise-Filtered Diverse Sampling, which filters outliers via deep-layer statistics and constructs frame-aware diverse calibration clusters to ensure stable quantization ranges. Comprehensive experiments demonstrate that QuantVGGT achieves the state-of-the-art results across different benchmarks and bit-width, surpassing the previous state-of-the-art generic quantization method with a great margin. We highlight that our 4-bit QuantVGGT can deliver a **3.7$ \times$** memory reduction and **2.5$\times$** acceleration in real-hardware inference, while preserving over **98\%** reconstruction accuracy of the full-precision counterparts. This demonstrates the vast advantages and practicality of QuantVGGT in resource-constrained scenarios.

Core Task Landscape

This paper addresses: **3D Reconstruction and Visual Geometry**

A total of **50 papers** were analyzed and organized into a taxonomy with **33 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Model Compression and Optimization**
- **Machine Learning Foundations and Paradigms**
- **Application Domains and Task-Specific Methods**
- **Research Methodology and Meta-Science**
- **Domain-Specific Empirical Studies**

Complete Taxonomy Tree

- 3D Reconstruction and Visual Geometry Survey Taxonomy
- Model Compression and Optimization
  - Post-Training Quantization for Vision Models ★ (1 papers)
  - [0] Quantized Visual Geometry Grounded Transformer (Anon et al., 2026) View paper
  - Test-Time Compute Scaling (1 papers)
  - [2] Scaling LLM Test-Time Compute Optimally can be More Effective than Scaling Model Parameters (Snell, 2024) View paper
- Machine Learning Foundations and Paradigms
  - Foundation Models and Transfer Learning (1 papers)
  - [48] Foundation Models -- A Panacea for Artificial Intelligence in Pathology? (Mulliqi, 2025) View paper
  - Multi-Task Learning Frameworks
  - Multi-Task Learning Theory and Algorithms (1 papers)
    - [34] A survey on multi-task learning (Yu Zhang, 2021) View paper
  - Multi-Task Learning in NLP (1 papers)
    - [37] A survey of multi-task learning in natural language processing: Regarding task relatedness and training methods (Zhang Zhihan, 2023) View paper
  - Multi-Task Learning in Vision and Multimodal Systems (2 papers)
    - [31] Touch the core: Exploring task dependence among hybrid targets for recommendation (Xing Tang, 2024) View paper
    - [33] MTRec: Multi-Task Learning over BERT for News Recommendation (Qiwei Bi, 2022) View paper
  - Training Objectives and Loss Functions (1 papers)
  - [20] Complement objective training (Chen, 2019) View paper
  - Quantum Machine Learning (1 papers)
  - [18] Is Quantum Advantage the Right Goal for Quantum Machine Learning? (Maria Schuld, 2022) View paper
- Application Domains and Task-Specific Methods
  - Embodied AI and Robotics (1 papers)
  - [13] A survey of embodied ai: From simulators to research tasks (Jiafei Duan, 2022) View paper
  - Vision and Language Understanding

    - Text Analysis and Representation (1 papers)
        - [11] Text algorithms in economics (Ash, 2023) View paper
    - Computer Vision Applications (2 papers)
        - [25] Measuring Robustness to Natural Distribution Shifts in Image Classification (Taori, 2020) View paper
        - [30] A survey of public datasets for computer vision tasks in precision agriculture (Yuzhen Lu, 2020) View paper
    - Natural Language Processing Shared Tasks (6 papers)
    - [21] SemEval-2022 Task 5: Multimedia Automatic Misogyny Identification (Elisabetta Fersini, 2022) View paper
    - [38] SemEval-2024 Task 8: Multidomain, Multimodel and Multilingual Machine-Generated Text Detection (Wang Yu-xia, 2024) View paper
    - [41] SemEval-2019 Task 5: Multilingual Detection of Hate Speech Against Immigrants and Women in Twitter (Valerio Basile, 2019) View paper
    - [42] SemEval-2020 Task 11: Detection of Propaganda Techniques in News Articles (Giovanni Da San Martino, 2020) View paper
    - [44] Predicting the Type and Target of Offensive Posts in Social Media (Zampieri, 2019) View paper
    - [46] The BEA-2019 Shared Task on Grammatical Error Correction (CHRISTOPHER BRYANT, 2019) View paper
    - Low-Resource and Multilingual NLP (1 papers)
    - [17] Participatory Research for Low-resourced Machine Translation: A Case Study in African Languages (Wilhelmina Nekoto, 2020) View paper
    - Software Engineering and Development (1 papers)
    - [40] From Idea to Implementation: Evaluating the Influence of Large Language Models in Software Development - An Opinion Paper (Yadav Sargam, 2025) View paper
    - Edge Computing and Task Offloading (1 papers)
    - [39] Task Partitioning and Offloading in DNN-Task Enabled Mobile Edge Computing Networks (Mingjin Gao, 2023) View paper
    - Semantic Communication Systems (1 papers)
    - [35] Deep Learning-Enabled Semantic Communication Systems With Task-Unaware Transmitter and Dynamic Data (Hongwei Zhang, 2022) View paper
    - Human-Computer Interaction and Gaming (1 papers)
    - [32] Aiming, pointing, steering: A core task analysis framework for gameplay (Bastian IlsÃ¸ Hougaard, 2024) View paper
- Research Methodology and Meta-Science
    - Literature Review Methods (3 papers)
    - [5] How to conduct systematic literature reviews in managementÂ research: a guide in 6 steps and 14 decisions (Seuring, 2023) View paper
    - [14] Metaâ□□analysis and traditional systematic literature reviewsâ□□What, why, when, where, and how? (J. Paul, 2022) View paper
    - [16] Geopolitical disruptions in global supply chains: a state-of-the-art literature review (Lukasz Bednarski, 2023) View paper
    - Research Design and Question Formulation (4 papers)
    - [15] A guide to research methodology: An overview of research problems, tasks and methods (Mukherjee, 2019) View paper
    - [22] Research problem (Kaushik Kumar, 2025) View paper
    - [28] Formulating a good research question: Pearls and pitfalls (Wilson Fandino, 2019) View paper
    - [50] Core research questions and methods (Christopher D. Mellinger, 2020) View paper
    - Research Paradigms and Epistemology (1 papers)
    - [6] Pragmatism as a research paradigm and its implications for social work research (Vibha Kaushik, 2019) View paper
    - Qualitative Research Methods (1 papers)
    - [8] Deductive Qualitative Analysis: Evaluating, Expanding, and Refining Theory (Stephen T. Fife, 2024) View paper
    - Bibliographic Infrastructure (1 papers)
    - [26] Web of Science (WoS) and Scopus: The Titans of Bibliographic Information in Today's Academic World (Raminta PranckutÄ□, 2021) View paper
    - Interdisciplinary and Decolonial Research (1 papers)
    - [49] Placing diverse knowledge systems at the core of transformative climate research (Ben Orlove, 2023) View paper
- Domain-Specific Empirical Studies
    - Biomedical and Clinical Research
    - Disease Classification and Diagnosis (2 papers)
        - [4] Updated classification of epileptic seizures: Position paper of the International League Against Epilepsy (Sandor Beniczky, 2025) View paper
        - [19] The Unrecognized Prevalence of Primary Aldosteronism (Jenifer M. Brown, 2020) View paper
    - Clinical Trials and Treatment Outcomes (3 papers)
        - [12] Management and outcome of primary CNS lymphoma in the modern era: An LOC network study. (Caroline Houillier, 2020) View paper
        - [29] Efficacy of Pembrolizumab in Patients With Noncolorectal High Microsatellite Instability/Mismatch Repair-Deficient Cancer: Results From the Phase II KEYNOTE-158 Study. (AurÃ©lien Marabelle, 2020) View paper
        - [47] Redefining the primary objective of phase I oncology trials (Mark J. Ratain, 2014) View paper
    - Epidemiological Cohort Studies (2 papers)
        - [9] Objectives, design and main findings until 2020 from the Rotterdam Study (M. Arfan Ikram, 2020) View paper
        - [45] Insomnia with objective short sleep duration and risk of incident cardiovascular disease and all-cause mortality: Sleep Heart Health Study (Suzanne M Bertisch, 2018) View paper
    - COVID-19 Impact Studies (1 papers)
        - [27] Impact of the COVID-19 pandemic on the core functions of primary care: will the cure be worse than the disease? A qualitative interview study in Flemish GPs (V. Verhoeven, 2020) View paper
    - Pharmaceutical and Chemical Research (2 papers)
    - [10] New Antibiotics for Multidrug-Resistant Bacterial Strains: Latest Research Developments and Future Perspectives (M. Terreni, 2021) View paper
    - [36] Research Progress on Antibacterial Activities and Mechanisms of Natural Alkaloids: A Review (Yumei Yan, 2021) View paper
    - Engineering and Industrial Applications
    - Quality Control and Manufacturing (1 papers)
        - [3] Quality Control to Reduce Appearance Defects at PT. Musical Instrument (Dikka Safriyanto, 2024) View paper

- Materials and Structural Engineering (1 papers)
  - [23] A review of the recent trends on core structures and impact response of sandwich panels (Quanjin Ma, 2021) View paper
- Maritime and Transportation Systems (1 papers)
  - [24] Towards the IMO's GHG goals: A critical overview of the perspectives and challenges of the main options for decarbonizing international shipping (Patrizia Serra, 2020) View paper
- Cybersecurity and IoT (2 papers)
- [7] CICIoT2023: A Real-Time Dataset and Benchmark for Large-Scale Attacks in IoT Environment (E. P. Neto, 2023) View paper
- [43] What Users Do besides Problem-Focused Coping When Facing IT Security Threats (H Liang, 2019) View paper
- Undefined or Placeholder Studies (1 papers)
- [1] The primary objectives of the show (Nairobi, 2022) View paper

Narrative

The field of 3D reconstruction and visual geometry encompasses a broad spectrum of research directions, organized into five major branches: Model Compression and Optimization, Machine Learning Foundations and Paradigms, Application Domains and Task-Specific Methods, Research Methodology and Meta-Science, and Domain-Specific Empirical Studies. The Model Compression and Optimization branch focuses on techniques to reduce computational and memory costs of vision models while preserving accuracy, including post-training quantization strategies that enable efficient deployment. Machine Learning Foundations explores core algorithmic paradigms such as multi-task learning frameworks and test-time compute scaling approaches like Scaling Test-Time Compute[2]. Application Domains spans diverse task-specific methods ranging from embodied AI systems to quality control in manufacturing, while Research Methodology addresses meta-scientific concerns including systematic review practices and research design frameworks. Domain-Specific Empirical Studies contribute datasets and findings from specialized areas such as medical imaging, agriculture, and cybersecurity.

Within the Model Compression and Optimization branch, post-training quantization for vision models represents a particularly active area addressing the challenge of deploying large-scale visual geometry systems under resource constraints. Quantized Visual Geometry[0] situates itself squarely in this compression-focused cluster, emphasizing efficient representation of geometric features through quantization techniques. This work contrasts with broader methodological studies like Systematic Literature Reviews[5] or Research Methodology Guide[15], which provide meta-level guidance on conducting research rather than proposing specific technical solutions. While application-oriented papers such as Quality Control Defects[3] demonstrate quantization benefits in industrial settings, Quantized Visual Geometry[0] appears more concerned with the fundamental compression mechanisms themselves, exploring trade-offs between model size, inference speed, and geometric reconstruction fidelity that are central to practical deployment of 3D vision systems.

Related Works in Same Category

No sibling papers were found in the same taxonomy leaf. A taxonomy-subtopic-level comparison will be produced instead.

Taxonomy-Level Summary

The original leaf focuses on post-training quantization techniques that compress vision models by reducing numerical precision after training is complete, aiming to reduce memory and computational costs. The sibling subtopic addresses test-time compute scaling, which instead allocates additional computational resources during inference to improve output quality. Both are inference-time interventions but pursue opposite goals: one reduces cost through compression, the other improves quality through increased computation.

Similarities: - Both operate at inference/test time rather than during training - Both aim to optimize the deployment and practical use of trained models - Neither involves retraining or modifying training procedures

Differences: - Post-training quantization reduces model size and inference cost through compression, while test-time compute scaling increases computational budget to enhance output quality - Quantization is a one-time model transformation applied after training, whereas compute scaling is a per-inference strategy - The original leaf is specific to vision and geometry models, while test-time compute scaling may apply more broadly across modalities - Quantization trades off some accuracy for efficiency, while compute scaling trades efficiency for improved accuracy

Suggested Search Directions: - Investigate whether post-training quantization and test-time compute scaling can be combined (e.g., quantized models with adaptive inference budgets) - Explore whether test-time compute scaling methods exist specifically for vision models that could complement or contrast with quantization approaches

Sibling Subtopics

- **Test-Time Compute Scaling** (leaves: 1, papers: 1)
- Scope: Strategies for improving model outputs by allocating additional inference-time computation.
- Exclude: Excludes parameter scaling and training-time methods; see Model Scaling and Architecture.

Contributions Analysis

Overall novelty summary. The paper introduces QuantVGGT, a post-training quantization framework specifically designed for billion-scale Visual Geometry Grounded Transformers. According to the taxonomy tree, this work resides in the 'Post-Training Quantization for Vision Models' leaf under 'Model Compression and Optimization'. Notably, this leaf contains only the original paper itself with no sibling papers, indicating a relatively sparse research direction within the broader 50-paper taxonomy. The work addresses unique challenges in quantizing large-scale 3D reconstruction models, particularly heavy-tailed activation distributions from special tokens and calibration instability from multi-view data.

The taxonomy reveals that the broader 'Model Compression and Optimization' branch contains only two leaves: the original paper's leaf and 'Test-Time Compute Scaling'. This suggests model compression for visual geometry models is an emerging area with limited prior work in the surveyed literature. Neighboring branches include 'Machine Learning Foundations' (covering foundation models and multi-task learning) and 'Application Domains' (spanning embodied AI to NLP shared tasks). The scope_note for the parent branch explicitly focuses on reducing computational costs while preserving performance, excluding training-time methods and architectural innovations, which helps contextualize this work's post-training focus.

Among the three contributions analyzed, the literature search examined 13 total candidates. The 'Dual-Smoothed Fine-Grained Quantization' contribution examined 4 candidates with 1 appearing to provide overlapping prior work, suggesting some precedent for smoothing-based quantization techniques. The 'Noise-Filtered Diverse Sampling' contribution examined only 1 candidate with no refutation, while the overarching 'First PTQ framework for VGGTs' claim examined 8 candidates with none providing clear refutation. These statistics reflect a limited search scope rather than exhaustive coverage, indicating the novelty assessment is based on top-K semantic matches within a constrained candidate pool.

Based on the limited 13-candidate search, the framework appears to occupy a relatively unexplored niche at the intersection of post-training quantization and large-scale visual geometry models. The sparse taxonomy leaf and absence of sibling papers suggest this specific application domain has received minimal attention in the surveyed literature. However, the partial overlap found for the smoothing technique indicates that while the overall framework may be novel, some underlying mechanisms draw from established quantization practices. The analysis covers semantic neighbors but cannot claim exhaustive field coverage.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Dual-Smoothed Fine-Grained Quantization (DSFQ)

Description: A quantization architecture that combines pre-global Hadamard rotation to disperse outliers and smooth heavy-tailed distributions with post-local channel smoothing to normalize channel-level variance. This dual-stage approach addresses the skewed activation distributions caused by data-independent special tokens in VGGT.

This contribution was assessed against **4 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. CLQ: Cross-Layer Guided Orthogonal-based Quantization for Diffusion Transformers

URL: View paper

Brief Assessment

CLQ Quantization[64] focuses on diffusion transformers for visual generation tasks, while DSFQ targets 3D reconstruction transformers (VGGT). The technical approaches differ: CLQ uses cross-block calibration and cross-layer parameter searching, whereas DSFQ addresses data-independent special tokens in VGGT through pre-global Hadamard rotation and post-local channel smoothing.

---

2. ConvRot: Rotation-Based Plug-and-Play 4-bit Quantization for Diffusion Transformers

URL: View paper

Brief Assessment

ConvRot[63] focuses on diffusion transformers for image generation using group-wise regular Hadamard transforms, while the original paper targets VGGT for 3D reconstruction with dual-stage smoothing (pre-global Hadamard rotation + post-local channel smoothing). The technical approaches and application domains differ substantially.

---

3. SmoothRot: Combining Channel-Wise Scaling and Rotation for Quantization-Friendly LLMs

URL: View paper

Prior Art Analysis

SmoothRot[62] demonstrates that the combination of channel-wise scaling and Hadamard rotation for quantization was proposed prior to the original paper's DSFQ contribution. Both papers address the same core problem (heavy-tailed distributions and outliers in transformer models) using the same dual-stage approach: Hadamard rotation to disperse outliers followed by channel-wise scaling to normalize variance. The technical mechanisms are nearly identical, with both applying rotation first, then scaling, to transform activations into quantization-friendly distributions.

Evidence

Evidence 1 - **Rationale**: Both papers describe the same dual-stage approach combining Hadamard rotation and channel-wise scaling to address outliers and heavy-tailed distributions, demonstrating that SmoothRot[62] proposed this combination prior to DSFQ. - **Original**: we propose a dual-stage smoothing scheme that globally disperses heavy-tailed distributions and locally balances channel variance, significantly reducing quantization errors - **Candidate**: smoothrot addresses the critical challenge of massive activation outliers, by integrating channel-wise scaling with hadamard transformations. our technique effectively transforms extreme outliers into quantization-friendly activations

Evidence 2 - **Rationale**: The original paper's DSFQ explicitly describes combining 'pre-global hadamard rotation and post-local channel smoothing,' which is the exact same technical approach as SmoothRot[62]'s integration of 'channel-wise scaling with hadamard transformations,' refuting the novelty claim. - **Original**: we introduce dual-smoothed fine-grained quantization, which integrates pre-global hadamard rotation and post-local channel smoothing to mitigate heavy-tailed distributions and inter-channel variance robustly - **Candidate**: smoothrot addresses the critical challenge of massive activation outliers, by integrating channel-wise scaling with hadamard transformations

---

4. RotateKV: Accurate and Robust 2-Bit KV Cache Quantization for LLMs via Outlier-Aware Adaptive Rotations

URL: View paper

Brief Assessment

RotateKV[61] focuses on KV cache quantization for LLMs using rotation techniques, while DSFQ addresses activation quantization for 3D reconstruction transformers (VGGT) with data-independent special tokens. The technical contexts and challenges are fundamentally different.

Contribution 2: Noise-Filtered Diverse Sampling (NFDS)

Description: A calibration dataset construction strategy that filters noisy outlier samples using deep-layer activation statistics and employs frame-aware clustering aligned with VGGT's inductive biases. This ensures a representative and stable calibration set for post-training quantization.

This contribution was assessed against **1 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Outliers and Calibration Sets have Diminishing Effect on Quantization of Modern LLMs

URL: View paper

Brief Assessment

Outliers Diminishing Effect[65] focuses on analyzing the diminishing role of calibration sets and outliers in modern LLMs, not on proposing calibration dataset construction strategies for vision-based 3D reconstruction models like VGGT.

Contribution 3: First PTQ framework for VGGTs (QuantVGGT)

Description: The first systematic post-training quantization framework specifically designed for Visual Geometry Grounded Transformers. It addresses unique challenges in quantizing billion-scale 3D reconstruction models through specialized techniques for handling data-independent tokens and multi-view data complexity.

This contribution was assessed against **8 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Pack-PTQ: Advancing Post-training Quantization of Neural Networks by Pack-wise Reconstruction

URL: View paper

Brief Assessment

Pack-PTQ[59] focuses on general neural networks for 2D image and 3D point cloud classification tasks, not specifically on Visual Geometry Grounded Transformers or 3D reconstruction models like VGGT.

---

2. APHQ-ViT: Post-Training Quantization with Average Perturbation Hessian Based Reconstruction for Vision Transformers

URL: View paper

Brief Assessment

APHQ-ViT[58] focuses on post-training quantization for Vision Transformers (ViTs) in general vision tasks, not on Visual Geometry Grounded Transformers (VGGTs) for 3D reconstruction. The candidate addresses different model architectures and application domains.

---

3. Empirical Research On Quantization For 3D Multi-Modal Vit Models

URL: View paper

Brief Assessment

3D Multi-Modal Quantization[56] focuses on quantizing 3D multi-modal ViT models, not Visual Geometry Grounded Transformers (VGGTs). The candidate's extremely limited context provides no evidence of addressing VGGT-specific challenges like data-independent tokens or multi-view calibration.

---

4. MGRQ: Post-Training Quantization For Vision Transformer With Mixed Granularity Reconstruction

URL: View paper

Brief Assessment

MGRQ Quantization[57] focuses on post-training quantization for Vision Transformers in 2D image classification tasks, not on Visual Geometry Grounded Transformers for 3D reconstruction. The technical domains and model architectures are fundamentally different.

---

5. Mesongs: Post-training compression of 3d gaussians via efficient attribute transformation

URL: View paper

Brief Assessment

Mesongs Compression[55] focuses on post-training compression of 3D Gaussian Splatting models for novel view synthesis, not Visual Geometry Grounded Transformers. The technical approaches differ fundamentally: Mesongs addresses Gaussian point attributes and spatial redundancy, while QuantVGGT targets billion-scale transformer quantization challenges.

---

6. 3dgstreaming: Spatial heterogeneity aware 3D gaussian splatting compression and streaming

URL: View paper

Brief Assessment

3D Gaussian Streaming[54] focuses on compression and streaming of 3D Gaussian splatting representations for rendering, not post-training quantization frameworks for Visual Geometry Grounded Transformers in 3D reconstruction tasks.

---

7. PTQAT: A Hybrid Parameter-Efficient Quantization Algorithm for 3D Perception Tasks

URL: View paper

Brief Assessment

PTQAT Quantization[52] focuses on hybrid PTQ/QAT for general 3D perception tasks (object detection, segmentation, occupancy) using CNNs and Transformers, not specifically on Visual Geometry Grounded Transformers or 3D reconstruction models like VGGT.

---

8. D3T: Dual-Domain Diffusion Transformer in Triplanar Latent Space for 3D Incomplete-View CT Reconstruction

URL: View paper

Brief Assessment

Dual-Domain Diffusion Transformer[51] focuses on 3D CT reconstruction using diffusion models in triplanar latent space, not post-training quantization frameworks for visual geometry grounded transformers.

Appendix: Text Similarity Detection

Textual similarity detection checked 15 papers and found 1 similarity segment(s) across 1 paper(s).

The following **1 paper(s)** were detected to have high textual similarity with the original paper. These may represent different versions of the same work, duplicate submissions, or papers with substantial textual overlap. Readers are advised to verify these relationships independently.

1. CLQ: Cross-Layer Guided Orthogonal-based Quantization for Diffusion Transformers

Detected in: Contribution: contribution_1

⚠ **Note**: This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

References

- [0] Quantized Visual Geometry Grounded Transformer View paper
- [1] The primary objectives of the show View paper
- [2] Scaling LLM Test-Time Compute Optimally can be More Effective than Scaling Model Parameters View paper
- [3] Quality Control to Reduce Appearance Defects at PT. Musical Instrument View paper
- [4] Updated classification of epileptic seizures: Position paper of the International League Against Epilepsy View paper
- [5] How to conduct systematic literature reviews in managementÂ research: a guide in 6 steps and 14 decisions View paper
- [6] Pragmatism as a research paradigm and its implications for social work research View paper
- [7] CICIoT2023: A Real-Time Dataset and Benchmark for Large-Scale Attacks in IoT Environment View paper
- [8] Deductive Qualitative Analysis: Evaluating, Expanding, and Refining Theory View paper
- [9] Objectives, design and main findings until 2020 from the Rotterdam Study View paper
- [10] New Antibiotics for Multidrug-Resistant Bacterial Strains: Latest Research Developments and Future Perspectives View paper
- [11] Text algorithms in economics View paper
- [12] Management and outcome of primary CNS lymphoma in the modern era: An LOC network study. View paper
- [13] A survey of embodied ai: From simulators to research tasks View paper
- [14] Metaâ□□analysis and traditional systematic literature reviewsâ□□What, why, when, where, and how? View paper
- [15] A guide to research methodology: An overview of research problems, tasks and methods View paper
- [16] Geopolitical disruptions in global supply chains: a state-of-the-art literature review View paper

- [17] Participatory Research for Low-resourced Machine Translation: A Case Study in African Languages View paper
- [18] Is Quantum Advantage the Right Goal for Quantum Machine Learning? View paper
- [19] The Unrecognized Prevalence of Primary Aldosteronism View paper
- [20] Complement objective training View paper
- [21] SemEval-2022 Task 5: Multimedia Automatic Misogyny Identification View paper
- [22] Research problem View paper
- [23] A review of the recent trends on core structures and impact response of sandwich panels View paper
- [24] Towards the IMO's GHG goals: A critical overview of the perspectives and challenges of the main options for decarbonizing international shipping View paper
- [25] Measuring Robustness to Natural Distribution Shifts in Image Classification View paper
- [26] Web of Science (WoS) and Scopus: The Titans of Bibliographic Information in Today's Academic World View paper
- [27] Impact of the COVID-19 pandemic on the core functions of primary care: will the cure be worse than the disease? A qualitative interview study in Flemish GPs View paper
- [28] Formulating a good research question: Pearls and pitfalls View paper
- [29] Efficacy of Pembrolizumab in Patients With Noncolorectal High Microsatellite Instability/Mismatch Repair-Deficient Cancer: Results From the Phase II KEYNOTE-158 Study. View paper
- [30] A survey of public datasets for computer vision tasks in precision agriculture View paper
- [31] Touch the core: Exploring task dependence among hybrid targets for recommendation View paper
- [32] Aiming, pointing, steering: A core task analysis framework for gameplay View paper
- [33] MTRec: Multi-Task Learning over BERT for News Recommendation View paper
- [34] A survey on multi-task learning View paper
- [35] Deep Learning-Enabled Semantic Communication Systems With Task-Unaware Transmitter and Dynamic Data View paper
- [36] Research Progress on Antibacterial Activities and Mechanisms of Natural Alkaloids: A Review View paper
- [37] A survey of multi-task learning in natural language processing: Regarding task relatedness and training methods View paper
- [38] SemEval-2024 Task 8: Multidomain, Multimodel and Multilingual Machine-Generated Text Detection View paper
- [39] Task Partitioning and Offloading in DNN-Task Enabled Mobile Edge Computing Networks View paper
- [40] From Idea to Implementation: Evaluating the Influence of Large Language Models in Software Development - An Opinion Paper View paper
- [41] SemEval-2019 Task 5: Multilingual Detection of Hate Speech Against Immigrants and Women in Twitter View paper
- [42] SemEval-2020 Task 11: Detection of Propaganda Techniques in News Articles View paper
- [43] What Users Do besides Problem-Focused Coping When Facing IT Security Threats View paper
- [44] Predicting the Type and Target of Offensive Posts in Social Media View paper
- [45] Insomnia with objective short sleep duration and risk of incident cardiovascular disease and all-cause mortality: Sleep Heart Health Study View paper
- [46] The BEA-2019 Shared Task on Grammatical Error Correction View paper
- [47] Redefining the primary objective of phase I oncology trials View paper
- [48] Foundation Models -- A Panacea for Artificial Intelligence in Pathology? View paper
- [49] Placing diverse knowledge systems at the core of transformative climate research View paper
- [50] Core research questions and methods View paper
- [51] D3T: Dual-Domain Diffusion Transformer in Triplanar Latent Space for 3D Incomplete-View CT Reconstruction View paper
- [52] PTQAT: A Hybrid Parameter-Efficient Quantization Algorithm for 3D Perception Tasks View paper
- [53] Point4Bit: Post Training 4-bit Quantization for Point Cloud 3D Detection View paper
- [54] 3dgstreaming: Spatial heterogeneity aware 3D gaussian splatting compression and streaming View paper
- [55] Mesongs: Post-training compression of 3d gaussians via efficient attribute transformation View paper
- [56] Empirical Research On Quantization For 3D Multi-Modal Vit Models View paper
- [57] MGRQ: Post-Training Quantization For Vision Transformer With Mixed Granularity Reconstruction View paper
- [58] APHQ-ViT: Post-Training Quantization with Average Perturbation Hessian Based Reconstruction for Vision Transformers View paper
- [59] Pack-PTQ: Advancing Post-training Quantization of Neural Networks by Pack-wise Reconstruction View paper
- [60] TinyVGGT: Lossless Post-Training Quantization for Visual Geometry Grounded Transformer View paper
- [61] RotateKV: Accurate and Robust 2-Bit KV Cache Quantization for LLMs via Outlier-Aware Adaptive Rotations View paper
- [62] SmoothRot: Combining Channel-Wise Scaling and Rotation for Quantization-Friendly LLMs View paper
- [63] ConvRot: Rotation-Based Plug-and-Play 4-bit Quantization for Diffusion Transformers View paper
- [64] CLQ: Cross-Layer Guided Orthogonal-based Quantization for Diffusion Transformers View paper
- [65] Outliers and Calibration Sets have Diminishing Effect on Quantization of Modern LLMs View paper